SURVEY: THE BRAIN

Who do you think you are?

Dec 19th 2006
From *The Economist* print edition

Modern neuroscience, says Geoffrey Carr, is groping towards the answer to the oldest question of all: who am I?

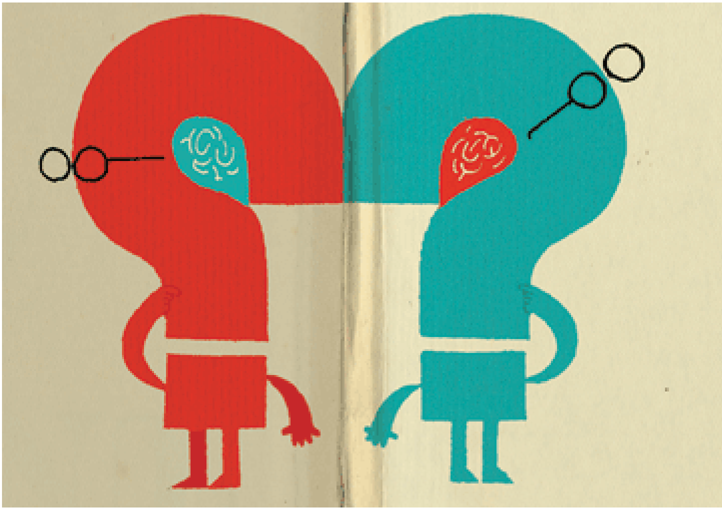

ON SEPTEMBER 13th 1848 a navvy called Phineas Gage was helping to build a railway in Vermont. As gang foreman, he had the job of setting explosive charges to blast a path through the hills near a town called Cavendish. While he was tamping down one of the charges with an iron bar, it went off prematurely, driving the bar clean through his head.

Accidents on construction projects happen all the time. The reason that people remember Gage's is that he survived it. Or, rather, his body survived it. For the Gage that returned to work was not the Gage who had stuck the tamping rod into that explosive-filled hole. Before, he had been a sober, industrious individual, well respected and destined for success. Afterwards, he was a foul-mouthed drunkard, a drifter and a failure. His identity had been changed in a specific way by specific damage to a specific part of his brain.

Gage's accident was intriguing because it cast light on the question of dualism. This is the idea that although the mind—the self—inhabits the brain, it nevertheless has an existence of its own and thus should not be equated with the brain. The sudden change Gage underwent suggested that brain and mind are not independent. If the essence of individuality can be changed by a physical accident, it implies that the brain is a mechanism which generates the self, rather than merely an organ which houses it. This observation moves the question “who am I?” from the realm of philosophy into the realm of science.

Thirteen years after the incident in Cavendish, a French neurologist called Paul Broca systematised the study of how brain damage affects the mind with the discovery that certain sorts of speech defect are the result of damage to part of the brain called the left temporal lobe (see article for a refresher course on brain anatomy and function). Local brain damage of this sort is known to neurologists as a lesion. Studying it therefore became known as the lesion method.

Broca's new method was taken up quickly. All sorts of strange neurological symptoms are now explained by specific brain damage. For example, an inability to perceive movement (even though the individual can see stationary objects) results from damage to part of the temporal lobe, and an inability to recognise faces is caused by damage to the fusiform gyrus. No one now questions the idea that particular parts of the brain specialise in particular activities.

Broca's revolution, though, is incomplete. On the face of things, its discoveries might have meant the end of dualism, but the world was not quite ready to embrace the mechanical explanation of self that the work of Broca and his successors implied. For much of the 20th century, a watered-down version of dualism based on the idea of the psyche prevailed. The distinction that psychiatry drew between neurological and psychiatric illness implied that there was a psyche (whisper not the word soul) that could somehow go wrong independently of physical symptoms in the brain.

When that idea was challenged by the effectiveness of physical drugs, such as antidepressants, in treating psychiatric illness, dualism returned in a different guise. Many people, most of whom would not regard themselves as dualists, think of the brain as being like a computer, and the mind as being like a piece of software that runs on that computer. But this analogy, too, is flawed. You do not have to do much damage to a computer to stop it being able to run programs. Yet as the case of Gage and numerous subsequent individuals has shown, the self can plod on, albeit changed, after quite radical brain damage.

The self in action

Broca's heirs, though, now have a range of new techniques with which to investigate the question. The best-known is a way of scanning the brain called functional magnetic-resonance imaging (fMRI). What makes it so powerful is that it records activity as well as anatomy. It can, if you like to put it that way, see the self in action. All you need to do is put someone inside an fMRI machine, give them a task to do and see which bits of the brain light up.

Naturally, the revolution in neuroscience brought about by this new technology has its critics. They point out that big conclusions are often drawn from small samples, that the changes in activity observed by fMRI are indirect (the technique measures blood flow and oxygen consumption rather than the electrical activity of nerve cells) and that the resolution is poor (individual points in an fMRI picture represent two or three cubic millimetres of brain tissue, which means hundreds of thousands of nerve cells). All these criticisms are justified. But these are early days. In science, time tells. The good studies are repeated and make the textbooks. The bad ones cannot be replicated and vanish down the memory hole.

Modern neuroscience has taken many directions, and this survey will not attempt to look at all of them. Instead, it will concentrate on four areas that may shed light on individual identity: the study of the emotions; the nature of memory; the ways that brains interact with each other; and the vexed question of what, exactly, consciouness is.

Such science is very much work in progress. Indeed, it is science of a type that would have been familiar to Broca and his contemporaries, for in many cases the researchers have only the haziest idea of where they are going. In the 19th century, when scientists were feeling their way towards big concepts such as the laws of thermodynamics, electromagnetics and the periodic table without really knowing what they were looking for, that was normal. These days there seem to be fewer new big concepts around, and experiments are often conducted in the expectation of particular results. But neuroscience is one area where big concepts certainly remain to be discovered. And when they are, they are likely to upend humanity's understanding of itself.

Captain Kirk's revenge

Dec 19th 2006
From The Economist print edition

Emotion is essential to human survival

ONE neuroscientist who could not be accused of dealing in small samples is Tor Wager, of Columbia University in New York. Dr Wager studies emotions—or, rather, he studies other people's studies of emotion. He has gathered together every fMRI study of emotion that he can lay his hands on—a total of some 150—and performed what statisticians call a meta-analysis. The result, illustrated below, is as close as anyone has yet come to producing an emotional map of the brain.

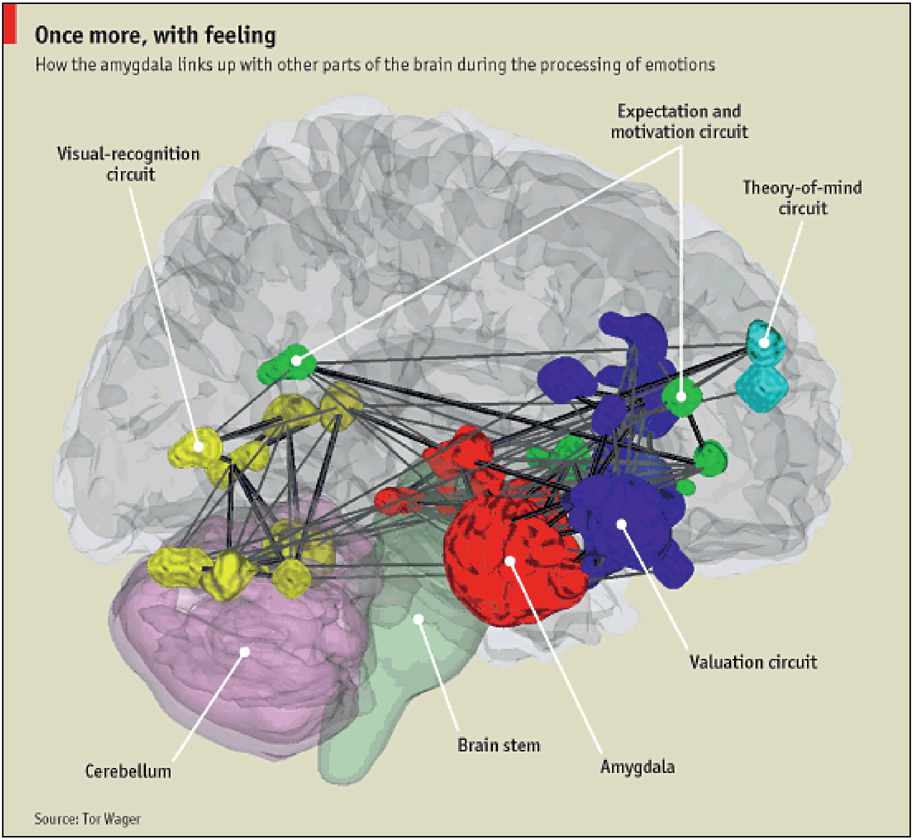

The experience of emotion is one of the most fundamental parts of an individual's identity. Most neuroscientists now recognise six basic emotions: anger, disgust, fear, joy, sadness and surprise. Dr Wager's map is a neat illustration of how fMRI can be used to see the links between different parts of the brain that are involved in a single process.

That people like Dr Wager can now study emotion scientifically shows how far things have come. For much of the 20th century, psychology sought to purge itself of the sin of anthropomorphism—that is, inappropriately ascribing human motives and feelings to other species. The tradition known as behaviourism approached animals as “black boxes”. Behaviourists stimulated them in different ways and recorded what happened. They did not ask what the animals felt. That both stymied comparative studies of emotion and put out of the scientific arena the question of how emotion evolved. Meanwhile anthropology, in a parallel ideological fit caused by the abuses of the eugenics movement, sought to expunge the idea that human behaviour had much in the way of a genetic underpinning. This was the infamous nature/nurture debate that lingers to this day.

Two people in particular came to the rescue: Paul Ekman and Joseph LeDoux. From the 1970s onwards, Dr Ekman, a psychologist at the University of California, San Francisco, challenged the anthropologists. He was responsible for the general agreement on the six basic emotions. He showed that the facial expressions associated with these emotions are universal, and therefore almost certainly plumbed in genetically.

In the 1980s Dr LeDoux, who is at New York University, challenged the behaviourists. Instead of rejecting anthropomorphism, he embraced it—though he did so carefully, noting the crucial importance of the word “inappropriately” in the ascription of human feelings to animals. He therefore studied fear, an emotion that no zoologist would doubt that mankind shares with other species, and used some of those other species to look inside the black box of the brain.

Now, as Dr Wager's ability to collect so many research papers suggests, studying emotion is all the rage. A glance at his map shows that many emotional pathways converge on two structures called the amygdalas. These are part of the limbic system, a collection of specialised structures in the middle of the brain, and it was Dr LeDoux who demonstrated their importance in a series of experiments carried out initally on rats. He used several techniques to confirm that the amygdalas are the most active part of the brain when the subject is afraid. He also produced fear by stimulating the neurons of the amygdalas with electricity. Subsequent work has shown that the amygdalas have the same role in people. Lose parts of them, as happens sometimes as a result of disease or surgery, and you may lose your ability to experience or recognise fear.

To start with, therefore, the amygdalas were thought of as the organs of fear. This, perhaps, is a good example of the sort of premature conclusion that critics worry about—because things turned out to be more complicated.

First, although the amygdalas do orchestrate fear, they seem to do so in the role of conductors as much as players. Certainly this emotional orchestra cannot play without the conductor, but the absence of the other instruments, whose functions are shown in Dr Wager's map, will also be noticed.

Second, the amygdalas also conduct other emotions. Since Dr LeDoux's pioneering work, further studies have linked anger, sadness and disgust with the amygdalas. They have also started to link other parts of the brain with particular emotions. Joy, for example, involves the amygdala's neighbour, the hypothalamus.

Genetics is starting to contribute to the study of emotion as well. The breakthrough came in 1993, with the discovery of a family (in the Netherlands, as it happened) that included an abnormally large number of violent criminals. The common factor in the criminal members of the family turned out to be the absence, due to a faulty gene, of an enzyme called monoamine oxidase A. This enzyme regulates a group of neurotransmitters that includes serotonin and dopamine. Serotonin- and dopamine-based neurons are both important for emotional responses.

At the time, the finding about monoamine oxidase A was widely reported as the discovery of “a gene for violence”. But violence is the expression of anger. Men without the gene were more easily angered. They had shorter fuses and were thus prone to spontaneous violent acts.

The Dutch study was followed up by one carried out in New Zealand by Terrie Moffitt, now of the Institute of Psychiatry in London. She took the nature/nurture question head on by demonstrating that the two interact, and in predictable ways. Again, the gene in question was the one for monoamine oxidase A. Like all genes, its activity is regulated by a DNA switch called a promoter. Monoamine-oxidase-A promoters come in two versions. Dr Moffitt found that a combination of one version and abuse during childhood really pushed people over the edge. The promoter alone, or abuse alone, resulted in some violent tendencies, but it was the mixture that made people really angry.

Illogical, captain

Humans share the basic emotions identified by Dr Ekman's work with other mammals. That helps to make them easy to study. But there is also a range of what are referred to, for want of a better phrase, as higher emotions. These are feelings thought to be confined, if not to humans alone, then to a small subset of large-brained mammals, several of whom are related to humans.

The list of higher emotions is not as well defined as that of the baser ones, but they include things such as guilt, embarrassment, shame and sympathy. What they have in common is that they depend not merely on what the person feeling them thinks about others, but on what the person feeling them thinks others are thinking about them. It is not the guilt or shame of the act itself, but the risk of being found out that provokes the emotion.

The evolution and function of these emotions is bound up with an area of research called theory of mind, to which this survey will return later. But, like basic emotions, the higher ones seem to have reliable neurological circuits whose location can be identified by fMRI.

Yoshiro Okubo, of Nippon Medical School in Japan, for example, has used fMRI to look at guilt and embarrassment. It is not easy to evoke such feelings in someone lying inside an MRI machine, but Dr Okubo thinks he has managed it. The results suggest that these emotions are handled in the medial prefrontal cortex (the middle of the front of the frontal lobe), the left posterior superior temporal sulcus (one of the furrows towards the side of the brain) and the visual cortex (towards the back of the brain).

It is surely no coincidence that much of the activity Dr Okubo found is in that characteristically human part of the brain, the enlarged cerebral cortex, rather than in the limbic system. And, as Dr Okubo points out, some of these areas are also associated with theory of mind.

The involvement of the frontal lobes is significant for another reason, though: it is the place where Phineas Gage took his hit. And that throws light on the question of what, exactly, emotions are for.

It is widely assumed that emotion and rationality are somehow opposed to each other, and that rational decisions are better than emotional ones. In fact, emotion and reason work closely together, as has been demonstrated by Antonio Damasio, the man who revived Gage's 19th-century fame in the 20th century.

Dr Damasio, who now works at the University of Southern California, is both a clinician and a researcher. He draws a parallel between Gage's case and those of some of his own patients. In particular, he has a patient called Elliot (in neuroscience, patients are often referred to by single names or initials to preserve their privacy) whose frontal lobe was damaged by a brain tumour. When the tumour was removed by surgeons, the damaged tissue was taken out too.

Like Gage, Elliot was a responsible individual with a good job (and in his case a family, too) before he suffered his brain damage. The outcome was somewhat different in that Elliot did not become a foul-mouthed wastrel; rather, he became obsessed with detail and stopped being able to make sensible decisions. The overall result was similar, though. He lost his job and his wife and ended up an outcast.

At first, Dr Damasio thought that Elliot's tumour had damaged his reason (both lesion studies and fMRI have shown that the frontal cortex is also the seat of the brain's reasoning powers). Tests, however, showed that what had gone instead were his emotions. Elliot no longer felt anything, and although he could summarise the choices available in a given situation as well as anyone else, without his emotions to guide him he could not actually make a choice. And, as probably happened with Gage, that loss of emotion also changed his self.

The survival value of things like fear, disgust and joy is obvious: run away from it; don't eat it; do more of it. But the idea that emotions shape all activity in adaptive ways is quite a subtle one. Rationality has its place. In the end, though, as fans of “Star Trek” will remember, it is Captain Kirk, the emotion-ridden human, not Mr Spock, the emotionless Vulcan, who has the nous to run the spaceship.

Brainbox

A history and geography of the brain

The most complex thing in the universe

A section through the brain. The temporal lobe is not visible in this diagram

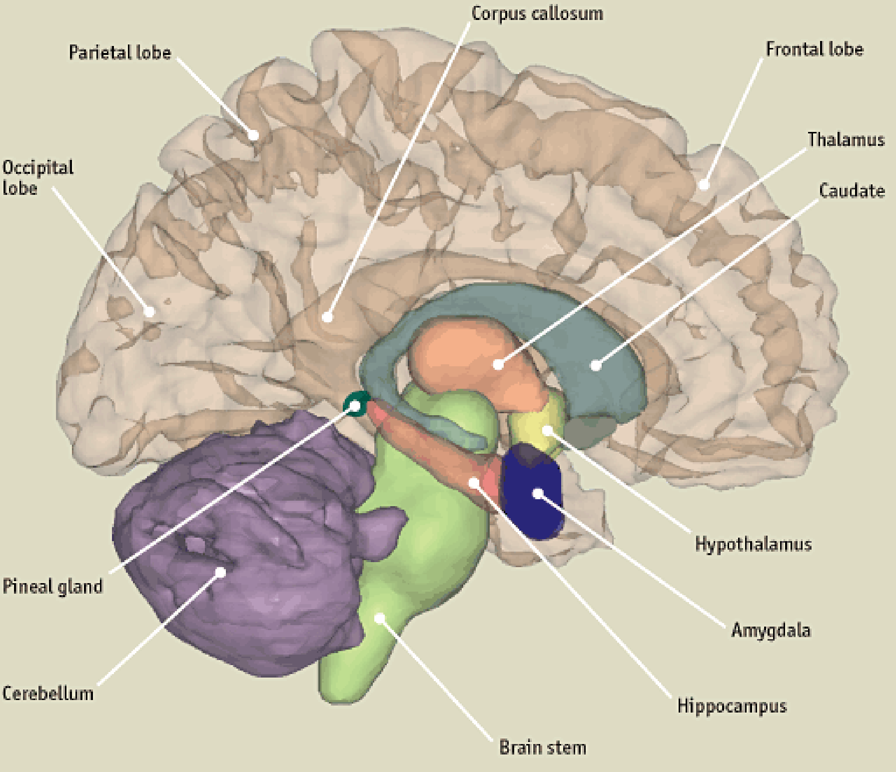

Source: Tor Wager

THE reason that people have brains is that they are worms. This is not a value judgment but a biological observation. Some animals, such as jellyfish and sea urchins, are radially symmetrical. Others are bilaterally symmetrical, which means they are long, thin and have heads.

Headless animals have no need for brains. But in those with a head the nerve cells responsible for it—and thus for sensing and feeding—tend to boss the others around. That still happens even when a long, thin animal evolves limbs and a skeleton. Bilateralism equals braininess.

A healthy human brain contains about 100 billion nerve cells. What makes nerve cells special is that they have long filamentary projections called axons and dendrites which carry information around in the form of electrical pulses. Dendrites carry signals into the cell. Axons carry signals to other cells. The junction between an axon and a dendrite is called a synapse.

Information is carried across synapses not by electrical pulses but by chemical messengers called neurotransmitters. One way of classifying nerve cells is by the neurotransmitters they employ. Workaday nerve cells use molecules called glutamic acid and gamma aminobutyric acid. More specialised cells use dopamine, serotonin, acetylcholine and a variety of other molecules. Dopamine cells, for example, are involved in the brain's reward systems, generating feelings of pleasure.

Many brain drugs, both therapeutic and recreational, work either by mimicking neurotransmitters or altering their activity. Heroin mimics a group of molecules called endogenous opioids. Nicotine mimics acetylcholine. Prozac promotes the activity of serotonin. And cocaine boosts the effect of dopamine, which is one reason why it is so addictive.

Apart from specialised nerve cells, there is a lot of anatomical specialisation in the brain itself. Three large structures stand out: the cerebrum, the cerebellum and the brain stem. In addition, there is a cluster of smaller structures in the middle. These are loosely grouped into the limbic system and the basal ganglia, although not everyone agrees what is what.

Most brain structures, reflecting the bilateral nature of brainy organisms, are paired. In particular, the cerebrum is divided into two hemispheres whose only direct connection is through three bundles of nerves, the most important of which is called the corpus callosum. (Many parts of the brain have obscure Latin names.)

This anatomical division of the brain reflects its evolutionary history. The brains of reptiles correspond more or less to the structures known in mammals as the brain stem and the cerebellum. In mammals, the brain stem is specialised for keeping the heart and lungs working. The cerebellum is for movement, posture and learning processes associated with these two things. It is the limbic system, basal ganglia and cerebrum that do the interesting stuff that distinguishes mammalian brains from those of their reptilian ancestors.

Soul-searching

The limbic system is itself divided. Some of the main parts are the hippocampus, the amygdala, the thalamus and the hypothalamus. The largest of the basal ganglia is the caudate. The pineal gland, which lies behind the limbic system, is the only brain structure that does not come in pairs. The 17th-century French philosopher René Descartes thought it was the seat of the human soul.

Descartes, however, was wrong. It is in fact the cerebrum's outer layer, the cerebral cortex, that is man's true distinguishing feature. The cerebral cortex forms 80% of the mass of a human brain, compared with 30% of a rat's. It is divided into lobes, four on each side. The rearmost one, called the occipital, handles vision. Then come the parietal and temporal lobes, which deal with the other senses and with movement. At the front, as you would expect, is the frontal lobe.

This is humanity's “killer app”, containing many of the cognitive functions associated with human-ness (although that most characteristic human function, language, is located in the temporal and parietal lobes, and only on one side, usually the left). Man's huge frontal lobes are the reason for the species' peculiarly shaped head. No wonder that in English-speaking countries the brainiest of the species are known as “highbrow”.

Dreamweavers

Dec 19th 2006
From The Economist print edition

The perfect memory is of everything and nothing

EVER since the unfortunate case of H.M., the subject of a lesion study second in fame only to that of Phineas Gage, neuroscientists have known that what the amygdala is to emotion, the hippocampus is to memory: if it is not the whole orchestra, then it is certainly the conductor.

H.M. lost the ability to form new memories when both of his hippocampuses were destroyed by radical surgery in the 1950s, and though he is now an old man, he still thinks of himself as the twenty-something who went into the operating theatre. He is palpably shocked whenever he sees himself in a mirror.

Memory is central to the question “who am I?” It is where the research is at its most nitty-gritty, with studies of the biochemistry of synapses and the action of individual genes, but also at its most esoteric, looking seriously at the function of dreams. It is, too, where cognitive neuroscience meets the videogames industry.

A few years ago Eleanor Maguire, of the Institute of Neurology in London, realised that her city has a resource available in no other large centre of population: taxi drivers who actually know where they are going. To become a London cabbie, you have to learn what is called the Knowledge: the location of every street within a 10km (six-mile) radius of the centre. This takes years of part-time study. If you fail the exam, you do not get a licence.

Dr Maguire and her colleagues used structural MRI (which predates fMRI, providing a static picture of anatomy) to study the hippocampuses of cabbies. They found that the shape of their subjects' hippocampuses varied with experience. As the Knowledge became consolidated, the back of the hippocampus seemed to grow while the front shrank.

With that result in their pocket, they have turned to videogames. One of the limitations of fMRI is that the scanners are heavy and unportable. But Dr Maguire wants to see what is going on in her taxi drivers' brains while they are using the Knowledge at work. So if the machine cannot travel around London, London has to travel round the machine. She makes this happen by employing a Sony videogame called “The Getaway” that can produce accurate representations of 110km (about 70 miles) of central London's roads. She uses it to project a lifelike image of driving round London onto a mirror visible from inside the machine. By this method, she is now studying how the hippocampus draws on other parts of the brain as her taxi drivers apply their Knowledge.

Knowledge or certainty?

Most researchers agree that long-term memory (as opposed to the short-term sort that can hold on to a telephone number long enough to dial it) comes in two varieties. One, known variously as explicit or declarative memory, records the salient details of an individual's life. For a taxi driver, the Knowledge is very salient. This form of memory involves the hippocampus. The other variety, implicit or procedural memory, involves the cerebellum and the basal ganglia. You may remember the anguish of individual violin lessons vividly via your hippocampus, but the finger movements required to play the instrument will be stored in your cerebellum. Even H.M. retains the ability to form new procedural memories, but his explicit memory has not grown since the time of his surgery.

To complicate matters further, there are two types of explicit memory. One, known as autobiographical or episodic memory, records the experiences themselves. The other, known as semantic memory, tries to generalise from these experiences. And there is evidence, to which Dr Maguire is trying to add, that the former is stored in the hippocampus whereas the latter is consolidated in the cerebral cortex.

One of the researchers trying to tease out the distinction between the two is Matthew Wilson, of the Picower Institute for Learning and Memory in Cambridge, Massachusetts. The Picower is one of several foundation-funded brain-research institutes that have sprung up in America over the past few years. (Jeffry Picower is a financier of biotechnology firms.)

Memory, as Dr Wilson observes, is like everything else in biology. It has evolved to serve a purpose and is honed for that purpose, which in this case is to react appropriately to the stimuli an animal meets in the environment by drawing on the experience of previous encounters. That is emphatically not the same as having a perfect memory for each of those encounters. Instead, memory should generalise from similar experiences and disregard the individual details. In other words, as time passes it should become more semantic and less autobiographical.

And indeed that is most people's everyday experience. The elderly are notorious for remembering every detail of their childhood but being unable to recall what they did last week. Such inability to remember details is often regarded as a failing, whereas so-called eidetic memory (or photographic memory, its more common name) is often admired by outsiders.

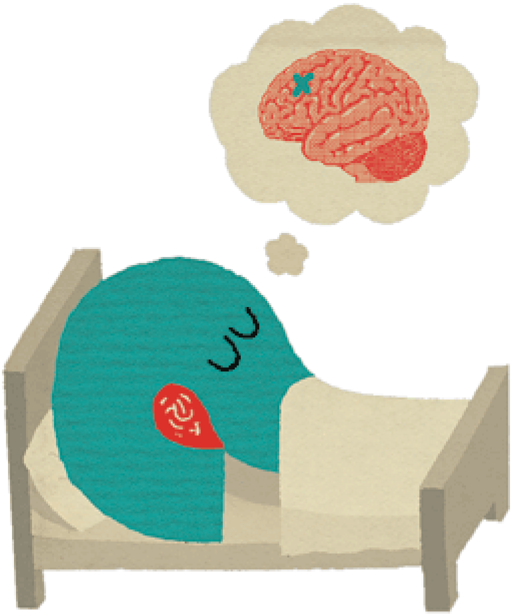

In Dr Wilson's view this perception is probably wrong. Indeed, an ideal memory would react like a behaviourist's black box. It would have generalised from experience to such an extent that individual events no longer need to be remembered at all; merely the appropriate response to the situation. So the fact that the elderly, who already have vast experience to draw on, do not waste precious storage capacity on adding things that will not aid their survival could well be the result of evolutionary adaptation rather than an indication of waning powers.

Dr Wilson studies memory formation by looking at rats. More specifically, he looks at rats dreaming—and day-dreaming—about what they have been up to. A connection between sleep, dreaming and the establishment of long-term memories has been known about for a while. Several years ago, he began recording the pattern of electrical activity in an animal's hippocampus as it learnt something about the environment, such as how to run round a particular maze, and showed that these patterns are recapitulated during what is known as rapid-eye-movement sleep, which in humans is the time for dreaming. This recapitulation seems to be crucial to memory formation.

He is now extending this work. He has shown that rats replay their experiences in their hippocampuses even when they are just resting, although, intriguingly, the pattern of electrical signals runs backwards at this time.

Learn in your sleep

Even more significantly, if electrodes are attached to neurons in the cortex that are connected to the hippocampus, part of the same pattern is seen there as well. However, there are differences between what is going on in the two places. When a rat is running a particular maze, the electrical pattern produced in the hippocampus is specific to that maze. Such patterns, though, share general features (similar corners in different mazes, for example, yield similar signals), and it is these general features that show up in the cortex. Dr Wilson interprets this as evidence of generalisation into semantic memory.

The recapitulation of experience in the form of neuronal firing patterns appears to be responsible for changing the pattern of synapses between nerve cells in ways that engrain particular memories by changing the way that information flows through the neuronal network. Dr Wilson's work does not explain exactly how those synaptic changes happen. But Elly Nedivi, one of his colleagues at the Picower, is one of those looking into the matter.

That long-term memory is encoded, at least in part, by changes in the strength of the synapses between nerve cells has been known since the pioneering work done by Eric Kandel of Columbia University in the 1960s and 1970s. Admittedly, Dr Kandel worked on a species of sea slug (an animal he chose because its neurons are easy to see and map), and at the time many of his contemporaries wondered whether the findings would hold true for more complex animals. But they did, and Dr Kandel's discovery of permanent changes in the strength and number of connections between neurons has become a cornerstone of the theory of memory.

Now, with the catalogues of genes provided by the Human Genome Project and its animal equivalents, it is possible to work out which genes are involved in these changes, and try to find out what they do. So far Dr Nedivi and her team have identified more than 360 genes that are unusually active in nerve cells during memory formation. They do this by looking for messenger molecules copied from genes in the nucleus and sent out to the protein-making apparatus in the cell body to tell it what to make. (It is these proteins that do the actual work in cells.) If particular messengers accumulate at synapses involved in memory, or seem to be associated with the growth of new axons and dendrites, that is a good indication that the proteins encoded by those messengers—and thus by the genes that generated them—have some role in the process.

The next stage is to find out what that role is. Using a fancy piece of genetic technology that can slice any given gene out of the nucleus and eliminate it, Dr Nedivi is doing what are, in effect, nano-lesion studies, starting with a gene that seems to be involved in making dendrites grow (dendrite growth failure is a cause of several forms of mental retardation). Her hope is that each gene's precise role can be worked out by seeing what happens in its absence.

Plato's cave

Working at the other end of the neurological scale—and across a large atrium from the Picower Institute—is Nancy Kanwisher of the McGovern Institute for Brain Research. Like the Picower, the McGovern is a privately supported autonomous satrapy of the Massachusetts Institute of Technology (Pat and Lore McGovern are entrepreneurs who made their money, respectively, in publishing and in computing). Here, Dr Kanwisher uses fMRI to look at where, exactly, in the brain various things are recognised.

Dr Kanwisher's discoveries go some way towards addressing the question of Platonic ideals—in other words, what is the essential property that makes an object, say, a table rather than a pile of firewood. What she has found is that certain pieces of cortex are able to extract these essential properties and thus react very strongly to particular sorts of objects. There is, for instance, the fusiform face area, which responds strongly (and only) to faces. The extrastriate body area responds similarly to images of human bodies or body parts. And the parahippocampal place area responds to images of places.

This specialisation makes sense. Faces, body parts and places are all important categories of natural object. Dr Kanwisher's latest discovery, though, is particularly intriguing. It is that there is also at least one area that handles a specific category of artificial object: written words. It, too, is always in the same place (a part of the cortex called the left fusiform gyrus). Somehow, all healthy developing brains not only work out that written words are a category to which it is worth allocating its own piece of neural anatomy, but find it easiest to accommodate that category in the same piece of wetware.

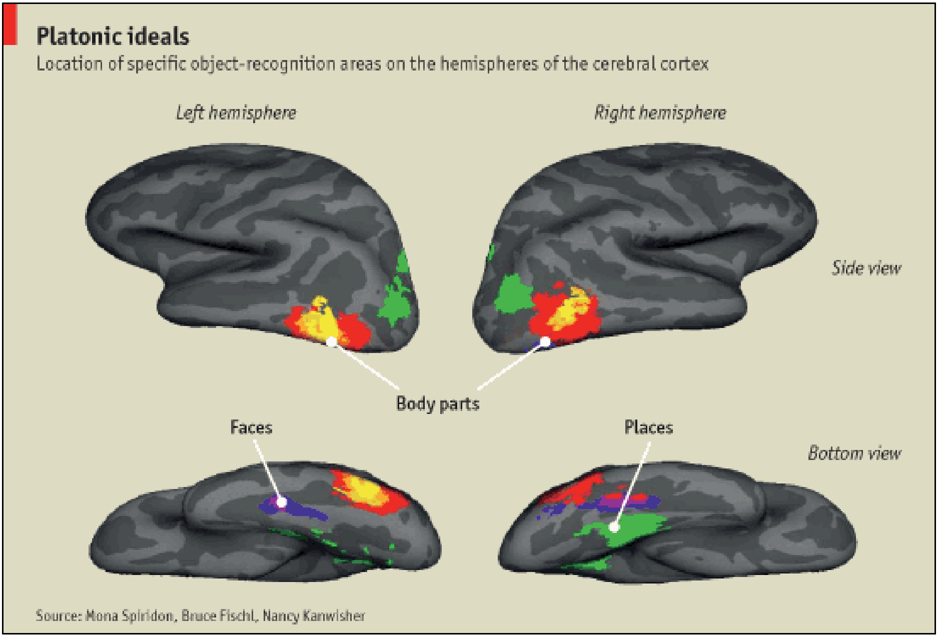

That could not have evolved specifically. Writing is probably too recent for natural selection to have done its work, and mass literacy certainly too recent. Understanding how such circuitry forms would yield an important insight into the logic of the mind. It might also indicate that the other specialised areas found by Dr Kanwisher are the result of developmental processes rather than evolutionary hard-wiring.

At the moment, it is hard to understand how the different circuits of neurons in the brain relate to one another. But Susumu Tonegawa, the head of the Picower, thinks he may be able to get closer to the answer by using gene elimination of the sort employed by Dr Nedivi to manipulate the circuitry of the whole brain. The tool that will allow him to do this is a project called the Allen brain map. It is named after Paul Allen, Bill Gates's partner in the founding of Microsoft. Mr Allen is even richer than Mr Picower and Mr and Mrs McGovern, and he, too, has paid for his own institute, which is based in Seattle.

Initially, the Allen Institute for Brain Science had but a single mission, which it completed in September. This was to create and publish a map of where in the brain particular genes are active, so that other researchers could use it in the knowledge that it was complete. Admittedly the brain in question is a mouse brain. But the genes of mice and men correspond closely, as does much of their neuro-anatomy, and Dr Tonegawa is not proposing to carry out his experiments on people. He does, however, plan to put the map to use in a rather clever way.

Many genes are actually switches that control the activity of other genes. By identifying particular switches that are active in only one part of the brain, he can co-opt those switches into activating his gene eliminator. Thus particular genes can be eliminated from one place without affecting others. That means he can shut down individual nervous pathways in the brain without affecting the others.

This technique, which he has only just begun to use, will take the lesion method to a new level of refinement. Although anatomists can trace connections between various parts of the brain by following the axons, finding our what the connections are actually for is much harder. If Dr Tonegawa's technique works, such discoveries will be easier to make. The workings of the whole of the brain, not just the parts concerned with memory, will be laid bare.

As others see us

Dec 19th 2006
From The Economist print edition

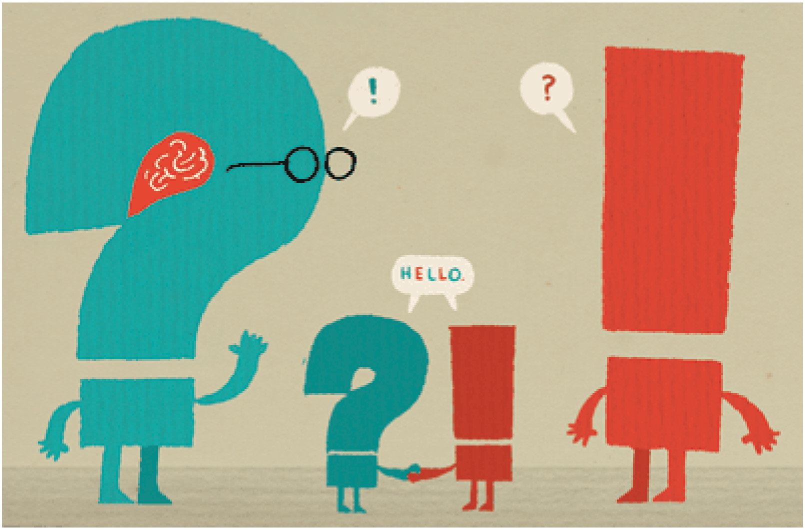

Dealing with people changes our minds

IN COUNTRIES where physical torture is illegal, the cruellest punishment that can be inflicted on an individual is solitary confinement. The “I” that exists by itself will surely go mad. Indeed, many students of the field think the evolutionary pressure that drove the enlargement of the human brain was not a need to survive in the natural environment but a need to negotiate the social one. The most obvious human attribute is language, and that is meaningless if there is no one to talk to. Less obvious, but just as important, is the ability to think oneself into the mind of another—in other words, to have a theory of mind. It is a combination of language and theory of mind that makes human society possible.

In science, as in other fields of endeavour, it helps to have a good slogan. “The language instinct”, coined by Steven Pinker, of Harvard University, is an excellent way of describing human powers of communication. But although Dr Pinker came up with the label, the idea of such an instinct was originally dreamed up by Noam Chomsky, who referred to it as “deep, universal grammar”. Ambitious claims about language have been made for other species, from parrots to dolphins to chimpanzees. None of these, however, has been shown to engage in the complexity of communication that people do.

Though Dr Pinker and Dr Chomsky disagree about the details—in particular on how the instinct evolved—there is a lot of behavioural evidence that the basic idea is right. The speed with which children learn the rules of speech is one piece of that evidence. It is hard to see how this could happen if what babies hear is not being plugged into some pre-programmed circuitry. Oddly, the difficulty of teaching the rules of writing is another piece of evidence. Writing is an artefact. Written language is no more complex than the spoken variety, but it is a recent invention and has not co-evolved with the language instinct. Children therefore struggle to master it. Perhaps the most persuasive behavioural evidence, though, is the way that the children of migrants in mixed-language communities (for example, sugar-producing islands in which slaves spoke different languages from each other and from their masters) are able to make up their own fully functional languages, known as creoles, in a single generation.

Besides the behavioural evidence, the mere existence of Broca's area, which governs speech production, and the speech-recognition area discovered a few years later by Carl Wernicke, points powerfully to the idea that a language instinct is hard-wired in by genetics. This is an area where fMRI has built convincingly on the original lesion studies. It has provided evidence that different parts of speech are dealt with by different parts of Broca's and Wernicke's areas. Not only are nouns and verbs processed in separate locations, but different types of noun may have their own areas, rather in the way that Dr Kanwisher's visual categories do. Such specialised locations have been claimed for categories such as animals and tools.

Nor is language processing merely a matter of decrypting and encrypting sound. Deaf people who communicate using sign languages (which have all the grammatical and syntactic features of spoken language) also do their processing in Broca's and Wernicke's areas. If they suffer damage to these areas, it shows up in exactly the same way that it does in those who can hear.

Taking the evidence in total, therefore, it seems likely that the Chomsky/Pinker theory is substantially correct. People have a specialised language instinct. The question is, why?

The Machiavellian mind

It is here that theory of mind—the ability not only to hypothesise what other minds are thinking, but to hypothesise what they are thinking about what you are thinking—enters the explanation. The evolutionary value of this is that people can anticipate the actions of others in a way that helps them. But with language, they can not only anticipate the actions of others, they can try to manipulate them. This idea was dubbed “the Machiavellian mind” by two of its originators, Andrew Whiten and Richard Byrne of St Andrews University in Scotland.

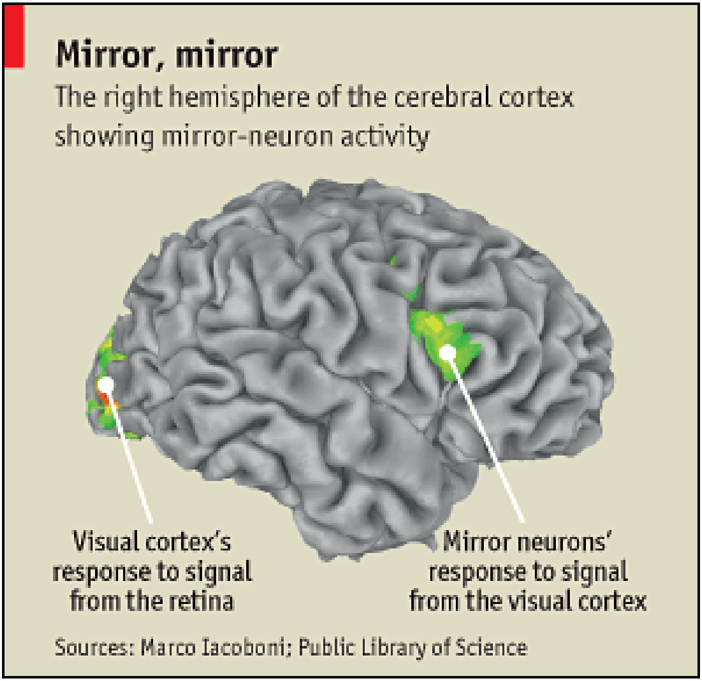

One of the most intriguing ideas about theory of mind comes from another sort of lesion study, an examination of the puzzling condition known as autism. People with autism find it hard to relate to their fellow human beings. Sometimes the condition is part of a wider range of problems, including low intelligence. But many people have autism pure and simple. Their intelligence is normal—indeed, it is often significantly above average—but their social relations are peculiar.

Simon Baron-Cohen, of Cambridge University, and Uta Frith, of University College London, think that autism results when an individual fails to develop a proper theory of mind. (Dr Baron-Cohen believes it is also an extreme manifestation of a general male tendency to be less “empathic” than females are.) All well and good, but modern neuroscience demands a mechanism. If one could be found, it might illuminate the way that theory of mind evolved in the first place.

Vilayanur Ramachandran, a researcher at the University of California, San Diego, thinks he may have found that mechanism: a failure of what are known as mirror neurons. A mirror neuron is one that is active both during the execution of a particular action or the production of a feeling by the individual concerned, and also when that individual observes the same action or feeling in another individual. In other words, it mirrors the actions and thoughts of others.

The first mirror neurons discovered (in the brains of monkeys) were concerned with predicting movement. However, mirrors of emotion have now been found as well, and it is these that interest Dr Ramachandran. He and his group showed, by looking at electrical signals, that the mirror-neuron system does not work properly in autistic children. The absence of relevant mirror neurons, he suggests, means the absence of theory of mind. For example, many autistic people have to learn the meaning of facial expressions by rote so that they can react appropriately.

That work has been followed up by Marco Iacoboni, of the University of California, Los Angeles, using fMRI. Dr Iacoboni discovered that the mirror neurons involved in grasping the intentions of others are in part of the right hemisphere that corresponds to Broca's area in the left hemisphere, and both also correspond with mirror-neuron sites in monkeys. Though no one has yet proved the case, it looks as though the evolution of language and the evolution of theory of mind might not only be two sides of the same coin, but might actually be different specialisations of the same basic structure.

I think, therefore I am, I think

Dec 19th 2006
From The Economist print edition

Consciousness awaits its Einstein

IN A building that looks, from the outside, like the villain's lair in an early James Bond film, a robot moves around. Called Darwin XI, it is the brainchild of Gerald Edelman. The building is the Neurosciences Institute in San Diego, California, and Dr Edelman is one of an eclectic group of researchers—some of them neurologists and some philosophers—who are trying to explain what is, perhaps, the biggest mystery of the human brain: the nature of consciousness. His approach is to build machines run by computer programs that work the way he thinks that brains work, and then see what happens.

Consciousness is the core of an individual's sense of self, yet, paradoxically, it is the most elusive concept in biology. Even framing the questions is difficult. Broadly, though, researchers have taken three approaches. One is the experimental method embraced by Dr Edelman. A second is to look for consciousness directly in the brain. The third is merely to sit and think about the question. Though empirical scientists sometimes scoff at it, this third method is not to be despised. After all, it was by sitting and thinking about some paradoxical results in physics that Albert Einstein was able to break out of the mental mould of classical physics and invent the non-commonsensical but scientifically successful theory of relativity.

Dr Edelman refers to his theory of consciousness as neural Darwinism. It combines two ideas. The first, as he charmingly puts it, is that “neurons which fire together, wire together.” This process of mutual reinforcement provides the selective pressure that is the prerequisite for any

Darwinian-based theory: to those neuronal networks that have shall be given, from those that have not, even what little they have shall be taken away. The resulting changes are the physical basis of learning.

The second part of Dr Edelman's theory is an idea he calls re-entrant mapping. The process of learning can be viewed as one by which reality (as perceived by the senses) is transformed into a representation of reality. Mathematically, this transformation is described as mapping. In Dr Edelman's model of the brain, however, there is a second process: the maps themselves are mapped by other groups of neurons. It is this phenomenon of different groups of neurons watching each other that he refers to as re-entrant mapping.

Whether neural Darwinism is truly a theory of consciousness is moot. It may not, for example, fully account for the feeling of actually experiencing things such as emotions that most people would think central to conscious experience. (Philosophers refer to such consciously experienced feelings as qualia.) As a theory of how brains work, though, it seems to have a lot going for it, for Dr Edelman has used it to construct a series of ever more complex robots that behave, in many ways, like animals. The latest, Darwin XI, has a range of senses: vision, hearing, touch and what Dr Edelman refers to as taste (but which is actually sensitivity to the electrical conductivity of what its “taste” organs are in contact with). It also has whiskers.

Darwin XI can do a lot. It can, for example, learn to navigate mazes in search of rewards, in the way that a laboratory rat does. It can develop preferences, thanks to a pleasure centre that generates what Dr Edelmann calls good taste in response to those rewards. And it can forget those preferences if they are no longer rewarding.

Other robots are able to perform similar tricks, but they have to be trained specifically to do so. The computer that runs Darwin XI can work things out for itself. It is loaded with virtual neurons, the initial strength of whose synapses with one another is allocated by a random number generator, and left to get on with things. It does have a bit of pre-ordained neuro-anatomy (in particular, it has been fitted with the equivalent of a hippocampus) but, like the local specialisation in a real cortex observed by people like Dr Kanwisher, most of the specialisation in Darwin XI simply emerges. This happens through the formation of specialised groups of neurons that resemble the specialised locations seen in real brains. The researchers know this because they can track changes in the way the virtual neurons connect to each other.

So is Darwin XI conscious? Well, it cannot speak, so no one can ask it. But the answer probably depends on whether you think a rat is conscious. That illustrates a big part of the problem of consciousness: no one can agree on who has it, let alone what it is. In fact, the questions are linked. There is a general feeling that what is special about humans is to do not with their being clever, but their being conscious in a different way from most other animals.

The inward eye

One feature of human consciousness that students of the field suggest might be unique is an awareness of self. The idea that self-awareness might be specific to humans and a few close relatives resulted from an experiment done three decades ago by Gordon Gallup, who now works at the University of Albany in New York state. This showed that chimpanzees (and, as subsequently emerged, other great apes) share with humans the ability to recognise themselves in a mirror, whereas monkeys and various other reasonably intelligent species, such as dogs, do not. A few species that are not apes have also passed the mirror test, including elephants and dolphins. But most animals fail it.

All the species that have passed have something in common: abnormally large cerebral cortices relative to the rest of their brains. Whether self-awareness simply emerges from a large cortex or whether selection for it necessarily results in one is unclear. Perhaps it is both. What is interesting about Dr Edelman's theory is that awareness of self is built into it. That, in essence, is what re-entrant mapping is.

Such self-awareness is not, however, indivisible. One treatment for serious epilepsy is to cut the corpus callosum and the other nervous connections between the two hemispheres of the brain, which stops the fit passing from one hemisphere to the other. This does not usually affect a person's everyday behaviour, but sometimes the two hemispheres have completely different personalities, and where that happens the individual's behaviour does change—indeed, he ceases to be an individual as the hemispheres fight for control of the body. The conflict often manifests itself in the person's hands, each controlled by a different hemisphere, trying to do opposing things. One hand may try to put on a piece of clothing, for example, while the other tries to remove it.

Tales of mystery and imagination

At first sight such cases seem extraordinary. But they are merely striking illustrations of a broader point: that in the brain nothing is ever quite what it seems, and experience and common sense are little use when formulating theories about the self. Two of the lesion studies mentioned in the introduction to this survey, dealing with the inability to perceive motion and recognise faces, arise from the fact that visual experience, which for those who can see is the dominant form of conscious experience, is a complete fabrication. What is consciously perceived is not a simple mapping of the images that fall on the retina. Instead, the signals from the optic nerves are deconstructed and re-formed in a process so demanding that it involves about a third of the cerebral cortex.

Even those with healthy brains get a hint of this in the form of optical illusions. These are patterns that the image-reconstruction process finds it confusing to deal with. An even more obvious discord between reality and perception is colour. The world is not really coloured, it just looks that

way because it is tremendously useful that it should, so the retina has cells that are particularly sensitive to three different wavelengths of light, and the brain weaves the signals from them together to create the phenomenon called colour.

Colours are good examples of qualia—the things that people feel that they are experiencing. Much of the philosophical side of the study of consciousness seeks either to explain qualia or to explain them away. They are, for example, at the heart of the question of dualism. For it is hard to ask what is generating them and what is perceiving them without concluding that the processes are separate.

Daniel Dennett, a doyen among philosophers of consciousness, disparagingly refers to the putative “observing self” in this scenario as a homunculus. He calls the mental stage on which the qualia supposedly act out their play the Cartesian theatre, after Descartes, the philosopher who thought the soul resided in the pineal gland. And he points out that exactly the same problem applies to how the homunculus would perceive its own qualia. Turn the theatre into a cinema, though, and Antonio Damasio quite likes the analogy. His twist is to place the observing self in the film itself, rather than in the audience. That is not a particularly easy idea to grasp, but it does seem to bear some relationship to Dr Edelman's idea of re-entrant mapping.

That something in the brain really is performing the role of an observing self is suggested by the work of Benjamin Libet at the University of California, San Francisco. Dr Libet used electroencephalography to look at brain activity during the process of making simple decisions such as when to move a finger. He showed that the process which leads to the act starts about three-tenths of a second before an individual is consciously aware of it. In other words, the observer is just that: an observer, not a decider. This may explain the feeling that most people have experienced at one time or another of having deliberately done something that they had not actually wanted or intended to.

Though Dr Libet's experiment is almost laughably simple, it pokes a stick in a very deep pond. A feeling of freedom to make conscious choices is at the heart of most people's sense of themselves. Even Freud, who popularised the idea of the unconscious, believed that conscious free-willed thought could override unconscious desires. One way of interpreting Dr Libet's work, though, could be that such free will is, like colour vision, simply a powerful illusion. An actor in a film, perhaps. But an actor reading from somebody else's script.

The truth, unsatisfactory though it is, is that no one really knows. Nor does anyone know where the next breakthrough will come from. Perhaps Dr Edelman, or one of his successors, will build a robot that can describe its own qualia-like experiences. Perhaps neuroanatomy will throw up a surprising, crucial observation. Or perhaps a bored, unregarded clerk will come to the rescue with an insight that dominates 21st-century thinking in the way that relativity dominated the 20th.